

Measuring and Fostering Diversity in Affective Computing Research

Isabelle Hupont , Songül Tolan, Pedro Frau, Lorenzo Porcaro , and Emilia Gómez

Abstract—This work presents a longitudinal study of diversity among the Affective Computing research community members. We explore several dimensions of diversity, including gender, geography, institutional types of affiliations and selected combinations of dimensions. We cover the last 10 years of the IEEE Transactions on Affective Computing (TAFFC) journal and the International Conference on Affective Computing and Intelligent Interaction (ACII), the primary sources of publications in Affective Computing. We also present an analysis of diversity among the members of the Association for the Advancement of Affective Computing (AAAC). Our findings reveal a “leaky pipeline” in the field, with a low –albeit slowly increasing over the years– representation of women. They also show that academic institutions clearly dominate publications, ahead of industry and governmental centres. In terms of geography, most publications come from the USA, contributions from Latin America or Africa being almost non-existent. Lastly, we find that diversity in the characteristics of researchers (gender and geographic location) influences diversity in the topics. To conclude, we analyse initiatives that have been undertaken in other AI-related research communities to foster diversity, and recommend a set of initiatives that could be applied to the Affective Computing field to increase diversity in its different facets. The diversity data collected in this work are publicly available, ensuring strict personal data protection and governance rules.

Index Terms—Diversity, ethics in affective computing, gender studies.

I. INTRODUCTION

THE Artificial Intelligence (AI) field has traditionally been dominated by a majority of white male researchers. This is why the last decade has shown the rise of initiatives and activities (often led by representatives of under-represented groups) directed at enhancing the diversity of the field by aiming at increasing the presence of women, researchers from different cultural origins or those with disabilities. The need for diversity in AI is

Manuscript received 16 June 2022; revised 16 January 2023; accepted 6 February 2023. Date of publication 10 February 2023; date of current version 1 March 2024. This work was partially supported by the European Commission under the HUMAIN project of the Joint Research Centre. Recommended for acceptance by Special Issue on Ethics in Affective Computing. (Isabelle Hupont and Songül Tolan contributed equally to this work.) (Corresponding author: Isabelle Hupont.)

Isabelle Hupont, Songül Tolan, Pedro Frau, and Emilia Gómez are with the Joint Research Centre, European Commission, 41092 Seville, Spain (e-mail: isabelle.HUPONT-TORRES@ec.europa.eu; songul.TOLAN@ec.europa.eu; pedro.frau-amar@ext.ec.europa.eu; emilia.GOMEZ-GUTIERREZ@ec.europa.eu).

Lorenzo Porcaro is with the Universitat Pompeu Fabra, 08002 Barcelona, Spain (e-mail: lorenzo.porcaro@upf.edu).

Digital Object Identifier 10.1109/TAFFC.2023.3244041

indeed acknowledged by both the academic community and policy makers. For instance, the *Ethical Guidelines for Trustworthy AI*, a document of the European Commission’s appointed expert group on AI, defines *diversity, non-discrimination and fairness* as one of the seven requirements that AI systems should fulfil in order to be trustworthy [1]. This requirement mentions the need to *avoid unfair bias, facilitate accessibility and universal design, and promote stakeholder participation*, ensuring respect to equality as a fundamental right expressed by the Universal Declaration of Human Rights (Article 7 [2]) and the European Union Charter of Fundamental Rights (Title III [3]).

The guidelines argue that, by fostering diversity at all stages of AI systems’ life-cycle, they could become more accessible to all and involve all relevant stakeholders. This includes the need for a diverse community of researchers and developers of AI systems. However, the current state of diversity in the AI field is yet to be assessed, and whilst several reports point out the crisis we are witnessing in terms of workforce diversity, e.g., [4], still shared practices for such assessment are lacking in the community. The urge of having robust indicators for characterizing participation in AI venues, and consequently for raising awareness of this crisis, is therefore the main reason motivating this work.

Affective computing is a research field related to the study and development of systems to recognize, describe, process and generate human affects [5]. The field deals at its core with AI human-centered applications, having a strong social and ethical impact, developing systems with a strong impact on humans such as emotion recognition and induction [6]. In order for these systems to be accessible to all, incorporate different views, and avoid biases and potential discrimination (e.g., with respect to gender, race, cultural background or languages), there is a need to ensure that research is carried out with a varied perspective, which should be reflected by the research community [7], [8].

The field is diverse *per se* in terms of disciplines, as it combines approaches from Computer Science, Cognitive Science and Psychology. A prominent example is Sentic Computing [9], a multidisciplinary multilingual approach for commonsense sentiment analysis that has induced many studies, notably on human rights and discrimination [10]. However, even inherent interdisciplinarity does not automatically imply diversity, as diversity covers multiple facets beyond interdisciplinarity [11]. In this work, diversity refers in the first place to the existence of variations of different characteristics in a group of people, more particularly in a research community. Evidence suggests that diverse teams outperform homogeneous groups on complex

tasks, including improved problem solving, increased innovation and more accurate predictions, all of which lead to better performance and results [12], [13], [14]. Diverse and inclusive scientific communities can generate new research questions not yet being asked in their particular discipline or culture, develop inclusive methodologies to better understand broader populations, and offer novel approaches to problem solving from multiple and different perspectives. Diverse groups have been shown to publish more articles, and these receive more citations per article [15]. Diversity thus enhances excellence, inclusion, generality and innovation.

The lack of tools to measure and monitor diversity is one of the limitations to assess the impact of efforts and policies. Hupont et al. [16] proposed a methodology to quantify diversity at the International Conference on Affective Computing and Intelligence Interaction (ACII), which considers gender, geographic location and institution type as diversity dimensions, and analyzes paper authors, keynote speakers and organizers. These results showed the limited diversity of the field, where the composition of researchers is persistently mostly men coming from Europe, Asia and North America. Some challenges of measuring diversity are the complexity of defining standard indicators, the lack of curated data (e.g., country, gender, institution type, topics), plus ethical concerns on the labeling of authors with gender information [17].

The current paper intends to complement previous literature, as discussed in Section II, by proposing a comprehensive methodological approach to measure, monitor and foster diversity in affective computing research (Section IV). For this purpose, we collect data on participants and papers of the three most prominent forums of the Affective Computing community: the IEEE Transactions on Affective Computing journal (TAFFC), the International Conference on Affective Computing and Intelligent Interaction (ACII), and the Association for the Advancement of Affective Computing (AAAC) (Section III). We use this data to first assess the diversity of the community's researchers in terms of gender (Section IV-A), institutional types of affiliations (Section IV-B), geography (Section IV-C) and the intersection of these dimensions. In a second stage, we assess how this diversity among people is related to the diversity of topics produced by them (Section IV-D). This is followed by an overview of existing diversity affinity groups and initiatives (Section V). Finally, we discuss how our findings translate into recommendations for initiatives and activities that could increase diversity in the Affective Computing community (Section VI). The diversity data collected and analysed in this study are publicly available at https://gitlab.com/humaint-ec_public/divinai-datasets.

II. RELATED WORK

The monitoring of diversity in scientific research communities is a practice that has been spreading in recent years. In the following, we present the most prominent studies and monitoring initiatives that have been undertaken to date in different scientific fields and, more particularly, related to Affective Computing. We discuss their focus, strengths and limitations, and compare them to this work.

A. Diversity Studies in Scientific Communities

The field of Neuroscience has been quite pioneering in diversity monitoring. The website *BiasWatchNeuro*¹ aims at raising awareness on the speaker composition of conferences, particularly with respect to gender representation. According to the current numbers in the website, almost half of studied conferences have ratios that fall below the base rate of women in the field [18]. The study presented in [19] centres on the number of black scientists being speakers at Neuroscience conferences, in particular comparing the situation before and after the death of George Floyd in 2020. Even though after 2020 the number of black speakers slightly increases, the under-representation of the black community is evident by the fact that in most analysed conferences there is no black speaker representation at all.

Notable works can also be found in the Geoscience field [20], [21], [22]. In particular, [20] analyses gender representation in 9 societies, 25 journals and 10 conferences evidencing the under-representation of women, especially in prestige roles such as conference organisers and journal editorial board members. Even though it is limited to a single edition of one conference, [21] discusses separately white versus black women participation, proving how the latter are much more under-represented than the former, who are already extremely under-represented in the field. Similar results have been obtained in disciplines such as STEM (Science, Technology, Engineering and Mathematics) [23] and Medicine (e.g., Paediatrics [24], Emergency Medicine [25] and Oncology [26]).

In the field of AI, some reports have looked at the composition of its members. The recent *2021 AI Index Report* produced by the Stanford University Human-Centered Artificial Intelligence Institute [27] evidences alarming results about the poor gender, ethnic, geographic and sexual orientation diversity in AI. More importantly, it highlights that the AI community currently lacks trustworthy means to collect information about the diversity of its members. Indeed, this report is based on a large-scale survey which clearly points –again– towards a field dominated by white men. On the positive side, it also shows that the *Black in AI* and *Women in ML* affinity groups (c.f. Table III) increased their number of participants in AI-related workshops in the last two years, which is indicative of the AI community starting to pay more attention to diversity and inclusion. The *AI Watch Index 2021* report by the Joint Research Centre of the European Commission [28] provides a good complementary view by reporting the diversity of participants, namely keynote speakers, authors and Program Committee (PC) members, in 5 top-tier AI scientific conferences in the period 2016-2020. Using the four Biodiversity-inspired indexes proposed by [17], [29] to measure gender, institutional and geographic diversity, the report evidences a slow but increasing diversity trend over the studied conferences.

B. Diversity Studies in Affective Computing

Diversity analyses of the Affective Computing research community are scarce. There are some recent bibliometric studies

1. *BiasWatchNeuro*: <https://biaswatchneuro.com/>

TABLE I

MAIN FOCUS, YEARS, METRICS AND DIMENSIONS ANALYSED IN STATE-OF-THE-ART DIVERSITY STUDIES. THE LAST ROW CORRESPONDS TO THE CURRENT WORK

MAIN FOCUS	YEARS	METRICS	DIVERSITY DIMENSIONS							
			Gender	Sex. orient.	Ethnicity	Age	Countries	Institutions	Topics	Cross
BiasWatchNeuro [18] Neuroscience: keynote speakers in >50 conferences and 4 journals	2015-2021	Women rate with respect to "base" rate	×							
Neuroscience: speakers in 18 conferences [19]	2019-2020	Percentages			×					
Geoscience: 9 societies, 25 journals and 10 conferences (organisation committee members) [20]	2016	Percentages	×			×		×		
Geoscience: speakers at 1 conference [21]	2017	Percentages	×		×				×	×
STEM: 1 society and 1 conference (speakers, attendees and poster presenters) [23]	2011-2015	Percentages	×						×	×
Medicine: speakers at 1 conference [25]	2016-2018	Percentages, speaking time	×							
AI Index Report [27] AI: survey data obtained from under-represented group members (women, queer, black) and participants in 1 workshop	2015-2020	Percentages	×	×	×		×	×		
AI Watch Index [28] AI: authors, keynote speakers and PC members in 5 top-tier conferences	2016-2020	4 diversity indexes	×				×	×		×
Affective Computing: authors, keynote speakers and PC members in ACII conference [16]	2005-2019	4 diversity indexes, percentages	×				×	×		×
This work Affective Computing: ACII conference (authors, keynote speakers and PC members), TAFFC journal and AAAC association	2011-2021	8 diversity indexes, percentages, clustering	×				×	×	×	×

analysing the past two decades of publications in the field [30], [31], [32] but, rather than on community members, they focus on identifying major publications, leading journals, key research topics, and most productive authors, institutions and countries in terms of number of published papers. They present nevertheless interesting findings directly or indirectly related to diversity as follows.

The growth rate of scientific production in the field differs slightly from one study to another, but is generally very high. According to Ho et al. [31], the annual growth rate of published papers is 12.5%. Pestana et al. [32] found that publications double every 4 years. Thus, Affective Computing is attracting more and more researchers every year which could potentially foster diversity in a natural way.

Additionally, results show that top contributing authors and institutions in the field come exclusively from academia and agree that most productive countries are USA, China, UK and Germany. Ho et al. [31] studied collaboration networks by country. Interestingly, they found two major collaborative networks: the "Asia Pacific cluster" with USA, China, Singapore and Japan, and the "European cluster" with Germany, UK and the Netherlands. While the diversity of the "Asia Pacific cluster" might bring some novel cross-cultural emotion studies, the "European cluster" risks falling into homophily. The presence of homophily in scientific collaboration is indeed a known phenomenon proven in several works [13], [33], showing how ethnicity, gender and affiliation factors are highly influential on shaping collaborations in the scientific community.

While bibliometric studies shed some light on the current landscape of the Affective Computing research community, they do not fully capture important facets of diversity such as the presence of under-represented groups (e.g., women, black people, other minority groups), institutions and countries. Our previous work presented in ACII 2021 [16] is, to the best of our

knowledge, the first attempt to quantify diversity in the Affective Computing community. It computes four diversity indexes for each ACII conference edition from 2005 to 2019, following the same methodology as the *AI Watch Index 2021* report [28], [29], and comparing them to those of top-tier AI conferences. Results show that ACII is well-positioned for gender and geographic diversity, though still far below the equality threshold for gender, and strongly dominated by Europe and North America in terms of countries. On the other hand, as expected from bibliographic studies of research, ACII lags far behind as regards institutional diversity, being totally dominated by academia (with imbalance ratios up to 1:15 for the presence of industry versus academia).

C. Towards a More Comprehensive Monitoring of Diversity: New Metrics and Dimensions

Table I summarises and compares the aforementioned state-of-the-art works on diversity. It shows their main focus (i.e., field and community members analysed), the years covered by each study, the metrics used to monitor diversity, the dimension(s) analysed (among: gender, sexual orientation, ethnicity, age, countries, institution types and topics) and whether cross-dimensional analyses are presented (e.g., gender \times country, gender \times topics, and so on). The last row of the table compiles the information corresponding to the present study.

As it can be seen, the gender dimension is considered by the vast majority of existing diversity works. However, the largest studies such as *BiasWatchNeuro* [18], [20] and [19] analyse the gender factor in isolation, without considering intersections between different dimensions. The most comprehensive work in terms of number of analysed dimensions is the *AI Index Report* [27], as information such as sexual orientation, ethnicity and age can only be collected by means of direct surveys

to community members. However, the latter does not present cross-dimensional views either.

The type of community members analysed varies among studies, although conference (keynote) speakers are the most common –and sometimes unique– target [18], [19], [21], [25]. Speakers have undoubtedly a key role in shaping the external prestige of a conference, but other actors such as PC members, authors and attendees are equally important when it comes to effectively evaluate diversity, as covered in few works [20], [23], [28] as well as in our previous study on ACII [16]. Interestingly, diversity among journal actors (authors, editorial boards) and societies/associations members is under-studied, especially in the AI field.

Finally, it is important to consider the duration (i.e., period covered) and the nature of the metrics used in these studies. Most works cover at most a period of 5 years and present simple percentages as quantitative evidence of diversity. While percentages are easy to understand by the wide public, there are some other interesting metrics that could be further explored. Traditional Biodiversity indexes (e.g., Shannon and Pielou) used in our previous work [16] are good examples, as they additionally allow for cross-dimension metrics. Other more sophisticated indicators applied to disciplines such as ecology [11] or telecommunications [34] could be also adapted to the analysis of diversity in a research community, but remain unexplored to date.

This work extends our previous study on the diversity of the Affective Computing community by including the last ten years of data from TAFFC journal’s authors and AAAC members, as well as new diversity dimensions (namely the topic dimension and cross-dimensional perspectives). It is important to highlight that our contributions go beyond the field of Affective Computing by adopting novel quantitative indicators of diversity among the members of a research community.

The monitoring of AI workforce diversification in an essential but not sufficient step in order to raise awareness and to promote diversity and inclusion initiatives, which may favor the advancement of racial [35] and gender literacy [36] in the AI community.

III. DATA

A. Sources and Nature of Data

To analyse the diversity of the Affective Computing research community throughout the past decade (January 2011 - December 2021), we compiled data from the three most prominent sources in the field:

- The IEEE Transactions on Affective Computing journal (TAFFC). We collected data from its papers and authors.
- The International Conference on Affective Computing and Intelligent Interaction (ACII). We compiled information from its papers, authors, keynote speakers and program committee members.
- The Association for the Advancement of Affective Computing (AAAC). We obtained data from the members of its executive committee.

For each identified member in the research community, be it a paper author, a member of a committee or any other role in the three aforementioned sources, we extracted the following information: gender (man or woman, inferred from names), country (inferred from the person’s affiliation country) and affiliation type. For affiliation types, we considered the eight categories in the Global Research Identifier Database (GRID²), namely: education, company, non-profit, healthcare, facility, government, archive and other. Additionally, we collected authors keywords from the papers published in TAFFC and ACII’s proceedings.

B. Data Collection Process

We followed similar though slightly different processes to extract all the data needed for the study, depending on the source.

The methodology followed for collecting data from ACII conference proceedings and the TAFFC journal is outlined in Fig. 1. The data collection pipeline is fully based on public domain information available online. We first gathered the list of papers published in both sources from January 2010 to December 2021 using the *Web of Science* (WoS³) global bibliographic database. In the case of ACII 2021, its proceedings were not yet indexed in WoS, thus we used the export tools from *IEEE Xplore*⁴ (publisher of ACII’s 2021 edition) instead.

For each of the 1175 so-collected papers (712 for ACII and 463 for TAFFC), we extracted the following information: title, publication year, list of authors (names) and their individual affiliations, plus the keywords provided by the authors when writing the paper. The final list of authors contains 2394 names for ACII and 1687 for TAFFC. As a second step, for each author we inferred gender from her/his name using the NamSor library⁵. From each affiliation, we obtained country and type by querying the GRID database, which currently contains about 102K research institutions. In order to ensure the quality and completeness of the data, we carried out some manual review and labelling of missing keywords (~ 7% of the total), affiliations (~ 3%), affiliation types (~ 67%) and countries (~ 2%). The high percentage of manual correction in features such as affiliation types (which was mostly due to double affiliations for certain authors, missing institutions in WoS or their different nomenclature in the GRID database) demonstrates the relevance of good quality data and the complexity of acquiring it, even if the papers are in principle available and indexed online.

In addition to ACII and TAFFC paper authors, we gathered and annotated fully manually the list of AAAC members, and ACII keynote speakers and organizers per edition. In the case of the AAAC, we got directly in touch with its current management board which kindly provided us the list (names) of the association’s executive committee members since its creation in 2007 for the purpose of this study. For ACII keynote speakers and Program Committee members, we obtained their names from conferences’ websites and proceedings’ foreword pages. Then,

2. GRID institutions database: <https://www.grid.ac/>

3. Web of Science: <https://clarivate.com/webofsciencegroup/solutions/web-of-science/>

4. IEEE Xplore Digital Library: <https://ieeexplore.ieee.org/>

5. NamSor library: <https://namsor.app/>

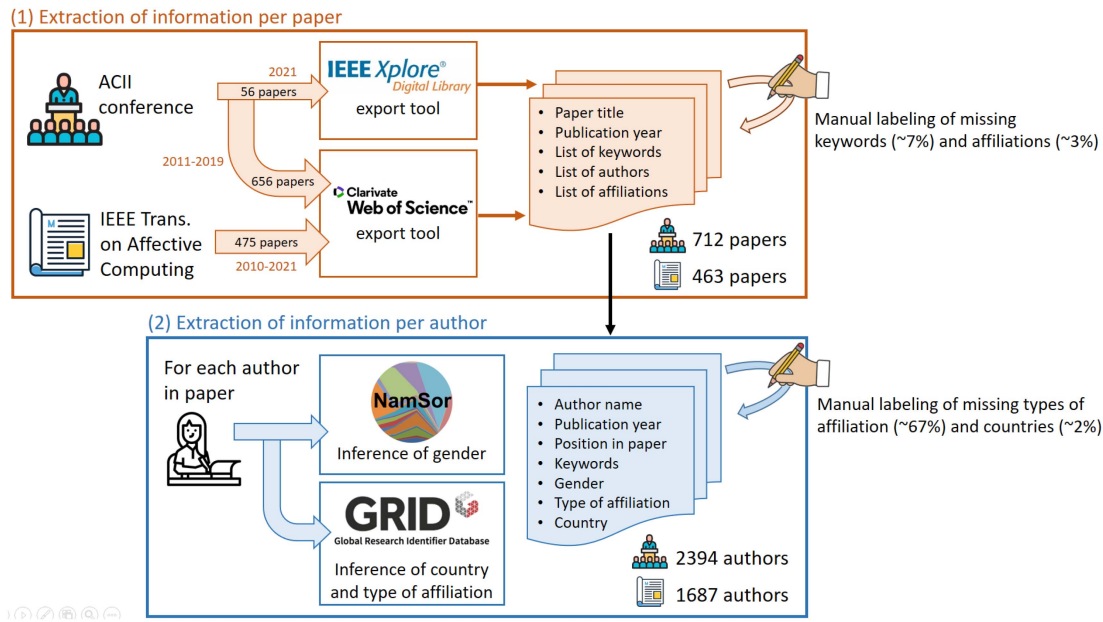


Fig. 1. Semi-automated process followed to collect per-paper and per-author data from ACII conferences and the IEEE Transactions on Affective Computing journal (years 2011 to 2021).

we manually labeled gender, country and type of institution (according to each person’s affiliation in a particular year and GRID classification).

C. Responsible Data Collection

In this section we discuss ethical considerations and risks related to the data collection performed in this study, as it deals with demographic data from the Affective Computing community members, namely their gender, and country and institution where they work(ed). Here we face a trade-off between, on the one hand, the need for diversity monitoring and reproducible research and, on the other hand, the risks of storing this personal information. There is indeed the need for further guidance on how to address the risks of demographic data collection for its responsible storage and use. McKane and Villeneuve [37] review related risks in the context of algorithmic fairness and provide a series of recommendations on how to feasibly collect, manage and employ demographic data. Although the present work does not focus on algorithmic fairness, we relate to it to discuss two of the key sensitive issues in our data collection process: privacy risks and labelling.

First, privacy risks require the implementation of technical methods to maintain individual non-identifiability throughout the data’s use. In this respect, we implemented a data protection procedure as follows. On the one hand, we stored raw data at an internal secure space, only accessible to the core research team and with the only purpose of ensuring full traceability of the results. On the other hand, we anonymised the data for an intermediate representation of members’ list to share with researchers and to comply with reproducibility practices in terms of indicator computations. The resulting anonymised data collected

and analysed during the current study are publicly available at https://gitlab.com/humaint-ec_public/divinai-datasets.

Second, labelling –and, more particularly, automatic labelling– procedures have associated several risks, such as the over-simplification of complex social concepts into categorical variables and potential mislabelling. For instance, our automated gender labelling process only allows for the classification of gender as binary classes.⁶ There have been ethical concerns regarding the use of systems for the automatic labelling of gender [38]. However, there is no study to date providing evidence on whether the quality of automatic data collection of gender is significantly different from manual labelling of such data. Similarly, we classified affiliations into one of the eight GRID categories, which might not fully capture the complexity of the matter (e.g., for persons with more than one affiliation). In this respect, we have carried out manual data correction to gender, institution and country information by checking authors web and social media presence for self-identification information. With these means, we have made our best efforts to minimize mislabelling.

IV. DIVERSITY INDICATORS

In general, the term “diversity” refers to the heterogeneity of elements in a set in relation to a class that takes different values, such as species in an eco-environment, or ethnicity in a population [39]. We measure diversity along four different dimensions: (1) gender, (2) institutional types, (3) geography, and (4) topics, where we also look at the intersection between gender and institutional type as well as gender and topic diversity. Since

⁶We refer to Wamesley (2022) for an online glossary on the diversity of gender identities: <https://www.npr.org/2021/06/02/996319297/gender-identity-pronouns-expression-guide-lgbtq>

the number of classes in the gender and institutional dimension is fixed across source and time, we assess gender and institutional diversity by simply computing the percentages of men/women and of each institutional type in the studied sources. In contrast, for every source the number of countries and topics can vary substantially over time. In order to account for this variation, we assess geographic and topic diversity using a three-dimensional measure of diversity. More precisely, this definition of diversity is derived from [11] who provides a cross-disciplinary general framework of diversity in science, technology and society. This work shows that diversity consists of three basic components:

- *Variety*: Refers to the number of classes in a set.
- *Balance*: Refers to the evenness of the distribution of elements across classes. Indicators of balance would be lower if a larger share of the elements are concentrated in only a few classes.
- *Disparity*: Refers to the degree of difference or distance between all classes.

Each of these components is necessary but individually insufficient for diversity. That is, we expect higher levels of diversity as classes increase, as the distribution of elements between classes become more even or if the difference (or distance) between classes becomes wider. Thus, for the geography and topic dimension we compute different metrics that each weight the different diversity components differently:

As metrics with a focus on balance, we compute the inverse of the Herfindahl-Hirschman Index (HHI), which is commonly used in economics to measure market concentration [40], and the Shannon Index (H') [41], which is commonly used in ecology to measure biodiversity [42]. HHI is computed as follows:

$$HHI = \sum_{i=1}^C p_i^2 \quad (1)$$

where $p_i = n_i/N$ is the proportion of elements from a given class i , i.e., the number of elements from this class n_i divided by the total number of elements N ; and C is the number of different classes. In a similar spirit, the Shannon index is computed as follows:

$$H' = - \sum_{i=1}^C p_i \ln p_i \quad (2)$$

As a metric with a focus on disparity, we compute the disparity metric (D) as follows:

$$D = \sum_{i=1, j=2}^C d_{ij}, i \neq j \quad (3)$$

where d_{ij} is the pairwise distance between classes i and j .

As a metric that combines all three diversity components providing a more general measure of diversity, we compute the Rao-Stirling metric (Δ) as proposed by [11]:

$$\Delta = \sum_{i=1, j=2}^C d_{ij} p_i p_j, i \neq j \quad (4)$$

Finally, in order to be able to compare results across diversity metrics, we compute the z-score of each diversity index across

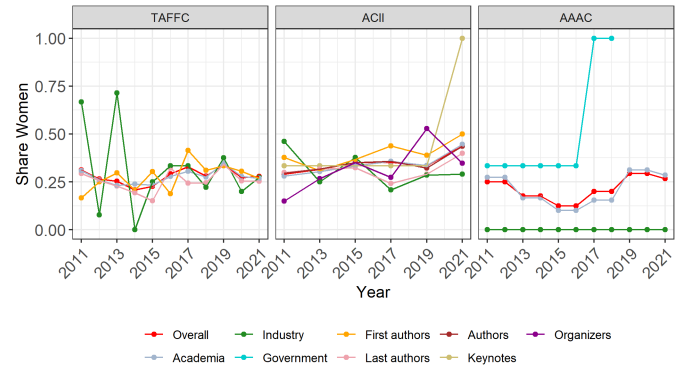


Fig. 2. Share of women over time by source and group.

all sources and years as follows:

$$z(div_{st}) = \frac{div_{st} - \bar{div}}{sd(div)} \quad (5)$$

Where div_{st} is the diversity index for source s and year t , \bar{div} is the mean of the diversity index across all sources and years and $sd(div)$ is the standard deviation of the diversity index across all sources and years.

As each of these diversity metrics considers each diversity component differently, we can interpret each metric from different perspectives. The balance-focused metrics Shannon and HHI increase as the distribution of elements across classes is not dominated by only one class. For instance, in the case of geographic diversity a balance-focused diversity index would be low if one outlet is mostly dominated by contributions from the USA and only a small share of the contributions come from different countries. The distance between these countries does not affect the balance-focused metrics. In contrast, the disparity metric would be lower for an outlet with contributions only from one continent than for an outlet with the same number of contributions but from different continents. The Rao-Stirling metric combines both of these diversity components but it would also have to be interpreted as a composite index of which the individual contributions of each diversity component are fuzzy. Finally, by allowing for different specifications of diversity, we ensure that our findings are robust across various specifications. We summarise the majority of computed indicators in Table II. In the following, we present the results for each studied diversity dimension separately.

A. Gender Diversity

We consider two different categories ($C = 2$) in the gender dimension: “women”, and “men”⁷. Since the number of categories is fixed over time, we indicate gender diversity through the share of women in each group. This comes with the benefit of simple interpretability of the result. We illustrate the share of women over time by group and source in Fig. 2.

Note first that for all three sources, the share of women overall always remains well below 40%. On average, AAAC

⁷We refer to Section III-C for a discussion of binary gender categories.

TABLE II

SUMMARY OF OBTAINED DIVERSITY METRICS BY SOURCE (TAFFC JOURNAL, ACII CONFERENCE AND AAAC ASSOCIATION), YEAR (2011-2021) AND DIVERSITY DIMENSION (GENDER, TYPE OF INSTITUTION, GEOGRAPHIC LOCATION AND TOPIC). NOTE THAT THE VALUES OF DIVERSITY INDEXES, NAMELY SHANNON (H'), HERFINDAHL-HIRSCHMAN (HHI), DISPARITY (D) AND RAO-STIRLING (Δ), ARE NORMALISED USING THEIR Z-SCORE FOR THE SAKE OF DIRECT COMPARISON ACROSS SOURCES AND YEARS

Source	Year	People	Papers	Share women	Institution diversity (share)			Geographic diversity				Topic diversity					
					Academia	Industry	Govern.	Countries	$z(H')$	$z(HHI)$	$z(D)$	$z(\Delta)$	Topics	$z(H')$	$z(HHI)$	$z(D)$	$z(\Delta)$
TAFFC	2011	63	18	0.31	0.86	0.05	0.03	14	0.28	0.49	-0.50	-0.86	10	-1.05	-0.41	-0.82	-0.34
TAFFC	2012	135	40	0.27	0.86	0.08	0.03	23	1.45	1.60	0.69	-0.37	19	-0.11	-0.12	-0.42	0.26
TAFFC	2013	116	37	0.25	0.89	0.06	0.00	20	0.41	0.17	0.13	-0.15	17	0.03	0.26	-0.53	0.48
TAFFC	2014	101	34	0.21	0.87	0.06	0.02	20	0.77	0.66	0.15	1.25	18	-0.10	0.31	-0.48	-0.25
TAFFC	2015	120	33	0.23	0.83	0.06	0.00	18	0.36	0.34	-0.01	1.73	16	-0.32	0.09	-0.58	-0.07
TAFFC	2016	147	32	0.29	0.88	0.02	0.03	18	-0.10	-0.51	-0.17	-1.14	22	0.50	0.46	-0.23	0.59
TAFFC	2017	138	41	0.33	0.90	0.02	0.00	24	0.62	-0.07	0.77	0.36	20	0.08	0.19	-0.36	0.10
TAFFC	2018	197	45	0.28	0.82	0.04	0.09	24	1.39	1.28	0.60	1.16	25	0.34	0.17	-0.02	0.39
TAFFC	2019	139	42	0.34	0.88	0.06	0.02	25	0.78	0.45	0.78	1.41	20	0.25	0.36	-0.36	0.52
TAFFC	2020	218	59	0.27	0.88	0.02	0.02	27	1.48	1.43	1.30	0.15	20	0.20	0.30	-0.36	0.44
TAFFC	2021	313	82	0.28	0.87	0.03	0.03	32	1.46	1.04	2.18	1.10	25	0.45	0.38	-0.02	0.50
ACII	2011	382	138	0.29	0.83	0.08	0.04	30	1.21	0.64	1.78	-0.95	40	1.05	0.66	1.50	0.42
ACII	2013	490	159	0.32	0.79	0.07	0.05	28	1.39	1.42	1.28	-1.05	43	1.17	0.62	1.89	0.57
ACII	2015	511	155	0.35	0.82	0.07	0.08	31	0.49	0.03	1.68	0.01	48	1.61	0.84	2.60	0.79
ACII	2017	331	96	0.35	0.87	0.06	0.01	22	-0.35	-2.04	0.53	0.16	7	-2.21	-0.71	-0.90	-3.14
ACII	2019	453	108	0.34	0.82	0.10	0.01	31	0.84	0.25	1.99	-0.53	32	0.87	0.58	0.60	0.59
ACII	2021	227	56	0.44	0.83	0.12	0.02	21	0.50	-0.05	0.20	-0.93	14	-1.91	-3.68	-0.68	-1.99
AAAC	2011	16		0.25	0.69	0.12	0.19	9	-0.71	-0.39	-0.82	0.65					
AAAC	2012	16		0.25	0.69	0.12	0.19	9	-0.71	-0.39	-0.82	0.65					
AAAC	2013	16		0.18	0.71	0.12	0.18	9	-0.64	-0.19	-0.82	0.45					
AAAC	2014	16		0.18	0.71	0.12	0.18	7	-1.38	-1.21	-0.90	0.48					
AAAC	2015	16		0.12	0.62	0.19	0.19	8	-0.85	-0.39	-0.90	0.33					
AAAC	2016	16		0.12	0.62	0.19	0.19	8	-0.85	-0.39	-0.90	0.33					
AAAC	2017	15		0.20	0.87	0.07	0.07	8	-0.76	-0.04	-0.90	-1.60					
AAAC	2018	15		0.20	0.87	0.07	0.07	7	-1.17	-0.60	-0.93	-1.22					
AAAC	2019	17		0.29	0.94	0.06	0.00	6	-1.89	-2.23	-0.96	-1.72					
AAAC	2020	17		0.29	0.94	0.06	0.00	6	-1.89	-2.23	-0.96	-1.72					
AAAC	2021	15		0.27	0.93	0.07	0.00	8	-0.82	-0.23	-0.90	-1.43					

has the lowest share of women, although there is a clear positive trend visible from 2016, which is not indicated by the other sources. Distinguishing between contributors from industry and academia reveals that the share of women is on average similar for TAFFC (although quite volatile until 2015 due to the smaller group size) but lower among industry contributors for ACII and AAAC. Interestingly, when considering author order for TAFFC and ACII, we find that the share of women among first authors is on average higher than among last authors. A reason for that could be that author order is usually correlated with seniority, where first authors are more likely to be junior researchers (PhD students, early postdocs) and last authors are more likely to be senior researchers (professors, supervisors). This is in line with the literature on “leaky pipelines” for women in science [43]. Finally, for ACII we distinguish between authors, keynote speakers and conference organizers but the differences between these groups in terms of share of women are not clear cut. In fact, except for 2019, the share of women among organizers is lower or equal to the share of women among authors. One notable point is that the keynote speakers of ACII 2021 were all women.

B. Institutional Diversity

We consider eight different categories in the institutional dimension ($C=8$, see Section III-A for the description of categories). Again, as the number of categories is fixed and limited over time, we indicate institutional diversity through the shares of each institutional type. The distribution of institutional types by source over time is illustrated in Fig. 3.

Naturally, as the data sources are all academic forums, the contribution from academia is largest across all sources and years. There are some interesting trends to note. While the contribution from industry for the TAFFC journal seems to

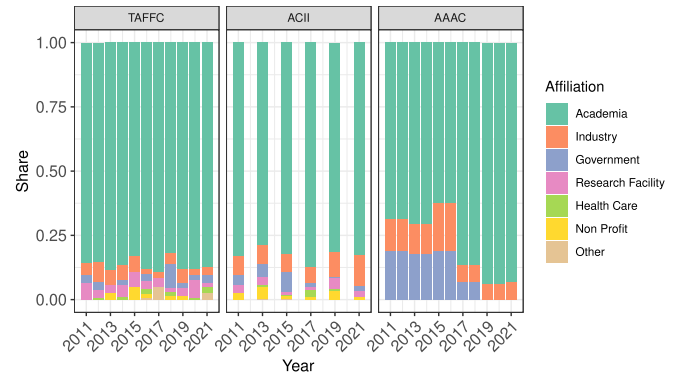


Fig. 3. Distribution of affiliation types over time by source.

decline over time, it seems to be increasing for the ACII conference. A reason for that could be that from the industry perspective conference participation comes with the additional benefit of being able to network and recruit potential employees with targeted skills. Furthermore, TAFFC seems to have more contributions from research facilities and healthcare than ACII. Another interesting development is the seeming disappearance of government contributions in all three forums. Explanations for these variations would have to rely on anecdotal evidence, which is why further research on the dynamics of institutional participation in Affective Computing research is needed.

C. Geographic Diversity

For the geographic dimension, classes are the affiliation countries of contributors, where the total number of classes varies over time. We compute diversity indexes for each source and each year. As measure of disparity, we compute the geographic

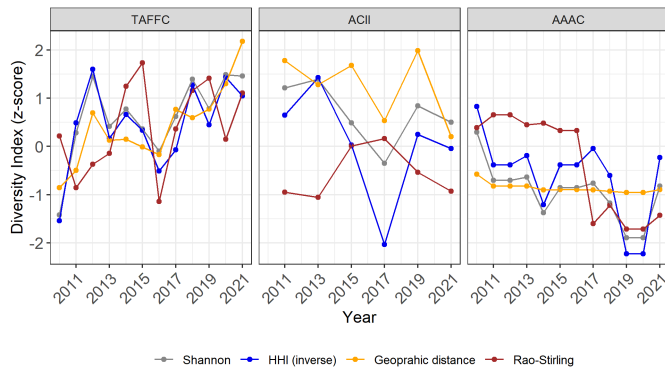


Fig. 4. Geographic diversity indexes over time by source.

distances between the capitals of each affiliation country and normalize this distance to a $[0,1]$ scale by dividing each distance by the maximum distance of each source-year group. Fig. 4 shows the z-scores of the geographic diversity indexes over time by source.

First, note that geographic diversity is quite volatile over time for all sources, which is probably caused by variation in the number of classes for every year (c.f. “countries” column of Table II). Moreover, the results suggest that geographic diversity is highly source-specific as the development of the indexes is notably different between the journal TAFFC, the conference ACII and the association AAAC. The main reason for this could be that the contribution to a conference like ACII requires travelling,⁸ which is subject to economic and location constraints of the contributor, while travelling is not necessarily required for contributions to a journal or association. Overall, it seems that geographic diversity is increasing for TAFFC, somewhat decreasing for AAAC and volatile for ACII, probably depending on the conference location. Interestingly, we find that the virtual ACII 2021 conference had lower than average diversity levels. A reason for that could be a lower motivation for participation in virtual events in the second year of the Covid-19 pandemic. Furthermore, while most diversity indicators are correlated, we find a notable difference between different types of diversity indexes for AAAC. While the disparity index remains relatively constant over time, the balance and rao-stirling indexes are decreasing. This suggests that the set of classes (affiliation countries) remains similar over time but the distribution of contributions across affiliation countries became less even. We explore this further in wordclouds of affiliation countries normalized by source in Fig. 5.

Indeed, Fig. 5 shows for TAFFC relatively clearly an increase in the number of authors (i.e., increase in font size) and countries. Moreover, while for 2010-2013 TAFFC the dominant contributor was the USA, this position is shared between USA and China in 2018-2021 with notable second-level European contributors from UK, France, India, Germany and Spain. As expected from the previous figure, the picture is less clear for ACII, especially in terms of number of authors and countries.

8.Note that ACII 2021 was virtual only due to the Covid-19 pandemic.

While the USA is dominating in terms of contributors for every year of the conference, there is indeed a variation in the other dominating countries depending on the location of the conference. For instance, ACII 2013 and 2019 took place in Europe, Geneva and UK, respectively, and many larger contributions came from European countries for these years. In contrast, ACII 2015 took place in Xian, China and for this year, China is a large contributor. In addition, the notably smaller and less diverse wordcloud in 2021 confirms the previous finding of low diversity for ACII 2021. Finally, the scarce wordclouds for AAAC underline previous findings, and, while from 2010-2017 AAAC still had members from Canada and China, in 2018-2021 AAAC consisted of members from the USA, Australia and some European countries.

D. Topic Diversity

We measure topic diversity using the same indexes as for geographic diversity. We identify topics based on a community detection approach on co-occurrences of author keywords. Note that we use author keywords (i.e., the keywords chosen by the authors themselves) as a proxy for topics rather than extracting them automatically from full paper texts using topic modeling techniques (we refer the reader to [44] for a recent overview on topic modeling algorithms). The reason is that we purposely want to capture the topics explicitly selected by the authors to identify their work with.

Links between keywords are weighted by the minimum conditional probability of their co-occurrence in a paper [45]. That is,

$$\omega_{ij} = \frac{k_{ij}}{\max(k_i, k_j)}, \quad (6)$$

where ω_{ij} is the weight of the link between keyword i and j , k_{ij} is the number of co-occurrences and k_i is the number of occurrences for keyword i . To make sure that the most often mentioned keywords do not form one big class, we remove links when both linked keywords are in the top 95th percentile of total occurrences. For all source-year groups, we run the Louvain community detection algorithm with the same specification [34]. With this method, node communities (i.e., topics) are greedily identified by comparing the density of connections within topics with the density of connections between topics, where the resulting number of topics is not pre-specified. The application domains of this community detection method ranges from Neuroscience [46] to Cybersecurity [47] as well as Social Science [45]. As topics we consider any identified keyword-community with more than two keywords. To obtain the paired distances between topics for the disparity index, we compute the inverse of the normalized paired similarities between topics. Paired topic similarity is the weighted sum of all links between all keywords from any two topics, where the weights correspond to ω_{ij} . We show the topic diversity indexes for TAFFC and ACII in Fig. 6.

First, note that all diversity indexes display very similar patterns, suggesting a high correlation between topic balance, disparity and variety. What is most striking about this result is the similarity with the geographic diversity indexes. Indeed, for both



Fig. 5. Author affiliation countries over time by source.

geographic and topic diversity, ACII depicts peaks for the years 2015 and 2019 and dips for 2017 and 2021. Similarly, we observe a positive trend in topic diversity for TAFFC, which is parallel to its geographic diversity. Thus, geographic and topic diversity are correlated. The degree of causality and the mechanisms of this relationship are yet to be explored in future research. However, two aspects suggest that this may not be merely a spurious correlation. First, ACII geographic and topic diversity depict high volatility in parallel. If there were a general trend over time due to an overarching policy towards more or less geographic and/or topic diversity, this pattern would depict a clearer trend.

Second, the number and diversity of topics is identified ex-post and agnostically from the set of already accepted papers. Unless the initial call for papers and the paper selection committee explicitly aimed for topic (instead of geographic) diversity, we can assume that geographic diversity pre-dates topic diversity in the production of research from these studied sources.

For a better understanding of the evolution of topics in detail, we present the largest clusters of co-occurring keywords (i.e., topics) with at least 10 keywords and 7 occurrences by time intervals separately for TAFFC and ACII in Fig. 7. To avoid that these graphs are dominated by only one generic topic,




Fig. 6. Topic diversity indicators over time by source.

we removed the keywords “affective computing”, “emotion(s)”, “emotion analysis”, and “emotion recognition” from this part of the analysis. Topics are labelled by the keyword with the highest degree in the cluster and the size represents the number of occurrences of the respective keyword. In line with [48] and [31] we represent these topics as four typologies that are structured by two parameters taken from network analysis. The first parameter (y -axis in Fig. 7), termed “degree”, represents the weighted number of links between the respective keyword and all other keywords in the network, where weights correspond to ω_{ij} . The degree can be seen as a measure for the maturity (or development) of the topic. The second parameter (x -axis in Fig. 7), termed “centrality”, measures the number of times the respective keyword lies on the shortest path in the network between two other keywords. The centrality can be seen as a measure for the relevance for the overall studied set of topics. Note that all indicated topics are relevant to the field to a certain degree, as the figure only shows larger topics (covering at least 10 keywords) and the most important keyword within that topic. Yet, we can still follow the characterization as proposed by [48]. That is, topics in the upper-right quadrant are well developed and relevant to the overall field, while topics in the lower left quadrant are not as well developed and less important to the overall field. The topics in the lower left quadrant could potentially be emerging or declining topics. Similarly, topics in the upper left quadrant are well developed but not so relevant to the overall field, suggesting that topics in this quadrant tend to be niche themes.

Note first about Fig. 7, that the most relevant and developed topics for TAFFC and ACII are changing over time. For TAFFC these are “affect analysis” in 2011-2013, “facial expression analysis/recognition” and “multimodal analysis” in 2014-2017, to “feature extraction” in 2018-2021. For ACII these are “physiological” in 2011-2013, to “multimodal” in 2014-2017, to “computational modeling” in 2018-2021. Interestingly, both of these developments suggest a potential trend in the focus on subfields (i.e., affects or physiology) to more methodological tools (i.e., feature extraction or computational modeling, which

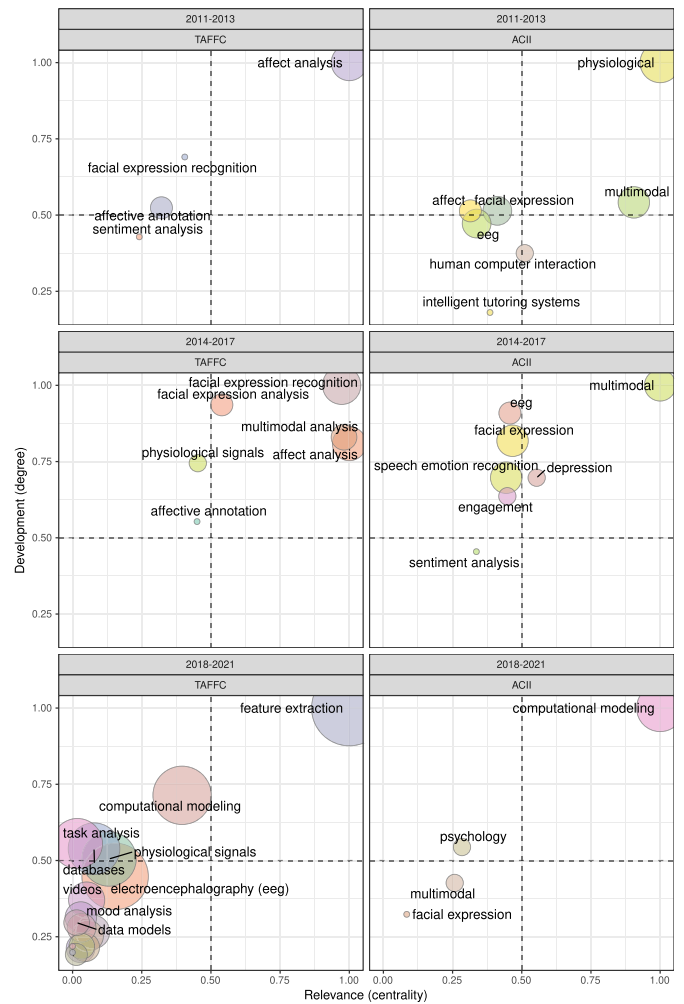


Fig. 7. Development of topics over time by source.

could be partly related to the strong rise of deep learning in AI). However, a deeper time-series analysis is necessary to confirm this hypothesis, as these types of cluster analyses for separate time periods do not take into account inter-temporal correlation.

Interestingly, we also confirm with Fig. 7 the somewhat opposing trends in topic diversity for TAFFC and ACII. While TAFFC starts in 2011-2013 with one important and few developed but still niche topics, ACII is seen at the beginning with two important topics and multiple less developed and less relevant topics. However, in time we see that other well developed topics emerge for TAFFC (facial expression and multimodal analysis is important in 2014-2017) and multiple new topics appear (e.g., task analysis, data models, videos, eeg), while ACII indicates over time a reduction in topics a stronger focus on one developed and relevant topics and fewer than TAFFC’s emerging/declining topics.

To address the question, whether women and men investigate different topics in the affective computing field, we conduct the same exercise of community detection on a network of keyword co-occurrences separately for papers with a woman as lead author and a man as lead author. We present the results, separately for




Fig. 8. Topics of papers with a woman versus a man as lead author.

TAFFC and ACII, in Fig. 8, where we also indicate the total number of papers (N) for each group.

First, note that the most relevant and developed topics of TAFFC and ACII are the same (or at least very similar) for papers with woman versus man lead authors. Naturally, we find differences in the total number of important topics as the share of women among first authors is much lower (133 versus 342 for TAFFC, and 270 versus 442 for ACII). However, we notice other interesting differences across woman-led versus man-led papers in both the topics and the characterization of topics in terms of relevance and development. More precisely, for TAFFC man-led papers focus mostly on one big topic, feature extraction, with many smaller topics in the area of emerging/declining themes. In contrast, a second well developed and relevant topic for woman-led papers is speech analysis, and two more well developed topics appear among the woman-led papers (physiological signals and eeg) that are not among the man-led papers. In contrast, for ACII man-led papers seem to indicate more well-developed and relevant topics (“physiological”, “multi-modal”, “affect”) than woman-led topics (“physiological”). In addition, there are some non-overlapping topics in the quadrant of well-developed topics that do not overlap, i.e. “computational modelling” and “facial expression” among man-led papers and “personality” among woman-led papers. At the same time, “deep learning”

appears as a more relevant (and almost well-developed) topic that does not appear among man-led papers. Exploring these topical differences requires a deeper bibliographic and time-series focused analysis, which goes beyond the scope of this paper. Nevertheless, these differences suggest that the gender of the lead author may have an impact on the topic of the paper.

V. DIVERSITY AFFINITY GROUPS AND INITIATIVES

From the previous sections, we can conclude that there is a wide room for diversity enhancement in the field of Affective Computing. In the following, we present a series of initiatives that have been successfully undertaken in other AI forums to enhance diversity, and that could potentially inspire our community, as will be discussed further in Section VI.

Table III provides a list of affinity groups actively involved in fostering diversity in the AI community. As it can be seen, most well-established groups put the focus on women in specific research areas including Machine Learning (*Women in ML - WiML*, active since 2007), Music Information Retrieval (*Women in MIR - WiMIR*, active since 2012), Recommender Systems (*Women in RecSys*, active since 2014) and Computer Vision (*Women in CV - WiCV*, active since 2015). More recently, since 2017, the focus has been widened to other demographics and under-represented groups such as Black people (*Black in AI*, created in 2017), Latin people (*LatinX in AI*, created in 2018), the LGBTQ+ community (*Queer in AI*, created in 2019), people with disabilities (*{Dis}Ability in AI*, created in 2019), Indigenous people (*Indigenous AI*, created in 2019) or African women (*African Women in AI*, created in 2022). We have performed an in-depth historical analysis of the diversity initiatives that have been promoted by both the organizers of international AI conferences and the aforementioned affinity groups. Table IV compiles the main findings, grouped by type of initiative.

The most traditional means of encouraging geographic diversity is to rotate the location of conferences between continents, to engage with different local communities, a measure also undertaken by ACII’s organisers. Although this facilitates the participation of researchers from different countries in alternate years, it does not help to improve the diversity of the conference in a given edition. In addition, involved geographic areas are generally Europe, Asia/Oceania and North America, leaving aside Africa and South America. To overcome this problem, financial support is needed, which has only very recently started to be granted by *Black in AI* and *LatinX in AI* in the form of conference fee and travel expenses payments. However, incentives to host conferences in Africa and South America are still lacking.

Specific workshops, dedicated panels and plenary sessions at conferences, and social events at main conferences are initiatives that have been successfully implemented since the early 2010s to make under-represented groups gain visibility. Following the trend in the emergence of affinity groups, they initially focused on women and were then extended in 2017 to other demographics and under-represented communities including Black, Latin, African, LGBTQ+ and people with disabilities. Recently, in

TABLE III
LIST OF DIVERSITY AFFINITY GROUPS IN THE ARTIFICIAL INTELLIGENCE RESEARCH COMMUNITY

AFFINITY GROUP	SINCE	FOCUS
Women in ML (WiML) ¹⁹	2007	Enhance the experience of women in ML, in order to help them succeed professionally and increase their impact in the community.
Women in MIR (WiMIR) ²⁰	2012	Promote the role of, and increase opportunities for, women, trans or non-binary at any career stage in the field of music information retrieval.
Women in RecSys ²¹	2014	Foster diversity and celebrate female role models in the recommender systems research community.
Women in CV (WiCV) ²²	2015	Foster the career and mitigate the isolation of female researchers working on computer vision.
Black in AI ²³	2017	Increasing the presence and inclusion of Black people in the field of AI.
Widening NLP (WiNLP) ²⁴	2017	Help to promote and support ideas and voices of under-represented groups in the Natural Language Processing community.
LatinX in AI ²⁵	2018	Latin professionals working on AI, ML and Data Science.
Queer in AI ²⁶	2019	People with diverse non-normative sexual orientations, romantic orientations and/or genders, corresponding to acronyms like LGBTQ+.
{Dis}Ability in AI ²⁷	2019	All those who experience barriers in accessing education due to having or being considered to have an impairment (e.g. physical or sensory impairments, people with learning difficulties, people with mental health or autism spectrum conditions).
Indigenous AI ²⁸	2019	Design and create AI from an ethical position that centers Indigenous concerns. The Indigenous term covers diverse communities in Aotearoa, Australia, North America and the Pacific.
African Women in AI (AWAI) ²⁹	2022	Promote knowledge sharing within the African women AI and ML community.

2021, the large and top-ranked AI conferences AAAI (Association for the Advancement of Artificial Intelligence Conference on AI⁹) and AAMAS (International Conference on Autonomous Agents and Multiagent Systems¹⁰) started to include an *open call for diversity and inclusion activities*, as a response to which affinity groups joined forces to propose a series of *Diversity workshops on AI*.¹¹

There are other fruitful –though less common– initiatives to help improving diversity in conferences. The role of *diversity and inclusion chair* was first created in the 2016 ISMIR Conference (International Society for Music Information Retrieval Conference¹²) so as to have a person devoted to take into account diversity in speakers and participants, as well as their inclusion and accessibility needs. Since 2018, some of the largest AI conferences such as NeurIPS, ICML, ICLR and RecSys also incorporate a *diversity and inclusion chair* in their organising committee. Nonetheless, ensuring an adequate level of diversity is not always easy, as there is a need for means to contact under-represented populations. Most affinity groups now provide directories,¹³ multimedia profiles of outstanding researchers¹⁴ (with interviews, podcasts, talks, etc.) and mailing lists¹⁵ that are particularly helpful, e.g., to invite a keynote speaker from a demographic minority or to make sure that a call for papers reaches an under-represented community. These directories also contribute to increase the visibility of these communities.

9. AAAI 2022 call for diversity and inclusion: <https://aaai.org/Conferences/AAAI-22/aaai-22-diversity-inclusion/>

10. AAMAS 2022 open call for diversity and inclusion: <https://aamas2022-conference.auckland.ac.nz/program/diversity-and-inclusion-activities/>

11. *Diversity workshop on AI* at AAAI 2022: <https://easychair.org/cfp/aidbei2022>

12. ISMIR 2021 organising team, including diversity chairs: <https://ismir2021.ismir.net/team/>

13. *LatinX in AI* directory: <https://lxai.app/PUBLIC-DIRECTORY>

14. WiML profiles: https://wimlworkshop.org/sh_projects/profiles/

15. *Queer in AI* mailing list: <https://groups.google.com/g/queerai>

Finally, beyond attendance and participation in conferences, other types of initiatives are oriented to further promote the scientific careers of individuals from minority groups. Mentoring allows early career professionals or research students, who are unaware of their options and lack role-models (e.g., due to local expertise limitations or financial hurdles) to get in touch with experienced senior industry or academics mentors and benefit from their experience. There are important mentoring programs ongoing since the early 2010s. Most of them debuted in the context of prestigious AI conferences (e.g., NeurIPS, ICML, AAAI, ICLR) and put a special focus on women. The Black¹⁶ and Latin¹⁷ communities have recently launched strong mentoring programs, independent from journals or conferences. The Latin community goes a step further by launching an official journal¹⁸ for community members to publish their latest research (e.g., works accepted in the workshops they organise).

VI. DISCUSSION

In this study we have analysed gender, geographical, institutional and topic diversity of the three most prominent forums of Affective Computing research, with the aim of raising awareness and fostering a discussion in the community on the need for more diversity in the field. Indeed, the main finding of this paper is that the state of diversity in Affective Computing could still be enhanced and there are already many active initiatives and affinity groups in the wider AI community that could help improve the situation.

With respect to the gender dimension, we find that the share of women remains below 40% across all three sources and most subgroups, such as institutional type or author leading position.

16. *African Women in AI* mentorship program: <http://wawiai.org/about-en.html>

17. *LatinX in AI* mentorship program: <https://www.latinxinai.org/mentorship-program>

18. Journal of LatinX in AI Research: <https://research.latinxinai.org/>

TABLE IV
SUMMARY OF DIVERSITY INITIATIVES THAT HAVE BEEN UNDERTAKEN IN INTERNATIONAL AI CONFERENCES. COLUMN YEAR SPECIFIES THE YEAR SINCE THE INITIATIVE HAS BEEN IMPLEMENTED

DIVERSITY INITIATIVE	FOCUS	ORGANISER(S)	CONFERENCE(S)	YEAR
ROTATION OF CONFERENCE LOCATION Rotate conference location to different geographical areas to improve participation of researchers in different countries.	The conference venue rotates between Europe, Asia/Oceania and America (mostly North)	Conference organisers	IJCAI, FG, ACII, ISMIR, FAccT	1969
WORKSHOP Specific workshop as a satellite event of a conference, sometimes with a different format (e.g. un-conference) and/or fully free access.	Women	WiML	NeurIPS, ICML	2008
	Women	WiCV	CVPR, ECCV	2015
	Black people	Black in AI	NeurIPS	2017
	Women and under-represented groups	WiMIR	ISMIR	2018
	LGBTQ+ people	Queer in AI	NAACL, FAccT, ICML	2019
	Latin people	LatinX in AI	CVPR, NAACL, ICML, ECCV, NeurIPS, ICCV	2021
	Women from Africa	AWAI	RecSys	2022
	Anyone who self-identifies within an under-represented demographic	WiNLP	AAAI, AAMAS	2021
DEDICATED PANEL AT MAIN CONFERENCE Panel at main conference, to outreach to the broader audience.	Women and under-represented groups	WiMIR	ISMIR	2015
		WiML	ICLR, AISTATS, CoRL, COLT	2020
SOCIAL EVENTS Social events at major conferences, e.g. breakfast or in-person meetups.	Women	WiML	NeurIPS, ICML, KDD, AAAI, COLT, CoRL, ICLR	2011
	Women	Women in RecSys	RecSys	2014
	Black people and women	Black in AI, WiML	NeurIPS	2020
	Latin people	LatinX in AI	ICLR, AAAI	2021
	Black people	Black in AI	ICML, NAACL	2021
	LGBTQ+ people	Queer in AI	ICLR, AAMAS	2022
CALL FOR DIVERSITY AND INCLUSION ACTIVITIES Open call for diversity and inclusion activities (e.g. workshops, roundtables, events) at a conference.	Under-represented groups such as women, LGBTQ+ persons and people of color	Conference organisers	AAAI, AAMAS	2021
DIVERSITY AND INCLUSION CHAIR Conferences incorporating a diversity, inclusion and/or accessibility chair to take into account diversity in speakers and participants.	Women	WiMIR	ISMIR	2016
	Inclusion, accessibility	Conference organisers, {Dis}Ability in AI	NeurIPS, ICLR, ICML	2018
		Conference organisers	RecSys	2021
DIRECTORY Directory of people from a minority community that provides visibility to under-represented groups. Particularly useful to look for invited speakers, area chairs, conference committee members, etc.	Women	WiML	-	2009
	Women and under-represented groups	WiMIR	-	2016
	Under-represented groups	WiNLP	-	2017
	Latin people	LatinX in AI	-	2018
	Indigenous people	Indigenous AI	-	2019
PROFILES Profiles of outstanding researchers in the minority group for visibility purposes, in form of interviews, podcasts, series of talks, etc.	Women	Women in RecSys	RecSys	2014
	Women and under-represented groups	WiML	NeurIPS, ICML, ICLR	2015
		WiMIR	ISMIR	2016
	Latin people	LatinX in AI	-	2018
MAILING LIST Mailing list within the community for networking, posting job positions, news, etc.	Women	WiML	-	2009
	Women and under-represented groups	WiMIR	-	2012
	Black people	Black in AI	-	2018
	LGBTQ+ people	Queer in AI	-	2021
FINANCIAL SUPPORT Journal fees, conference, summerschool registration and travel costs covered, and family support.	Women	WiML	NeurIPS, ICML	2015
	Women	WiCV	CVPR, ECCV	2015
	Women and under-represented groups	WiMIR	ISMIR	2016
	Black people	Black in AI	AAAI, CVPR, ACL, ICLR, COLING, FAccT	2017
	Latin people	LatinX in AI	Any related to AI	2018
	LGBTQ+ people	Queer in AI	Any related to AI	2019
	Women, under-represented groups and attendees with disabilities.	ICML, {Dis}Ability in AI	ICML	2022
	Women	WiML	NeurIPS, ICML, AAAI, KDD, AISTATS, COLT, ICLR	2009
MENTORING PROGRAM Mentoring program for networking and support (e.g. discuss career aspirations, receive scientific orientation on a research topic).	Women	WiCV	CVPR, ECCV	2015
	Women, LGBTQ+ and other under-represented groups	WiMIR	ISMIR	2016
	Black people	Black in AI	-	2017
	Latin people	LatinX in AI	-	2018
	Women from Africa	AWAI	-	2022
SPECIFIC JOURNAL Specific journal for community members, e.g. in the form of an archival publication of accepted papers in workshops.	Latin people	LatinX in AI	Journal of LatinX in AI Research	2018

And, while the trend appears positive, it is still rather slow. In addition, we find a lower share of women among contributors from industry, compared to academia, and among last authors compared to first authors, which suggests the presence of a “leaky pipeline” in affective computing careers.

Regarding institutional types, our results show that Affective Computing research is, naturally, well-dominated by academia, where contributions from industry are increasing, especially for the ACII conference, and government contributions are declining. More research is needed to understand drivers of these trends.

At the geography dimension, we find that geographic diversity is increasing for TAFFC, location-dependent for ACII and somewhat declining for AAAC. However, all three sources show that the overwhelming majority of Affective Computing research still comes from the USA, which is closely followed by China and with contributions from European countries in the second place. Contributions from Latin America or Africa are almost non-existent.

Lastly, one of the most interesting findings is that diversity in the characteristics of researchers seems to affect the diversity of topics. Indeed, geographic and topic diversity follow highly similar patterns over time, and paper topics are on average different depending on the gender of the lead author.

To address the identified shortcomings in the diversity of researcher characteristics and potentially obtain the benefit of more topic diversity, we formulate from our analysis the following recommendations:

- *Define diversity targets and monitoring tools.* The clear definition of diversity targets through indicators that can be monitored can help direct efforts and diversity funds, and potentially accelerate positive diversity trends.
- *Ensure data quality for monitoring.* As mentioned in [37], self-identification in demographic data linked to diversity incurs in least risk, as requiring data subject participation ensures that they have more awareness and control over their own data. However, this awareness can raise concerns about responsible data use. Thus, our recommendation is to ask authors for their demographic data (e.g., at conference registration time, when submitting a paper to TAFFC, through surveys) while informing them about their contribution to better diversity monitoring. High-quality data would allow for an accurate and reproducible monitoring effort, and informed consent would bring agency to the community in the diversity efforts and related monitoring process.
- *Assess the suitability of the diversity initiatives* reviewed in this study in the context of Affective Computing. It is important to assess, for each initiative, the suitability for the Affective Computing community in terms of needs, motivation, resources and impact. For instance, it may be adequate to see which community members may be motivated to lead certain initiatives (e.g., social activities, mentoring programs or affinity groups), the financial resources available (e.g., for travel grants) and be strategic in selecting those providing more impact and requiring less resources. Multi-faceted diversity indicators that are based on high quality data can help guide these decisions.

- *Strengthen relationships with existing affinity groups in AI.* We have identified key affinity groups from under-represented minorities that are very active in fostering diversity in AI, with excellent outcomes in the last 5 years. Getting in touch with them, having access to their mailing lists to promote Affective Computing outlets or inviting them to co-organise workshops are actions easy to implement in the short term and that could attract these minorities into the field. Particularly, we have seen the very scarce presence of researchers from Africa and Latin America, thus *Black in AI*, *Black Woman in AI* and *LatinX in AI* could be interesting targets. We have also discussed the lack of information regarding the sexual orientation and gender self-identification of our community members, which is an issue that could be potentially addressed in tandem with *Queer in AI*. Finally, *{Dis}Ability in AI* could help to identify and break barriers that Affective Computing members might have been experiencing.

VII. CONCLUSION

This work builds on Hupont et al. [16], presented at ACII 2021. Here, we extend the original study in terms of data basis (adding TAFFC and AAAC data), facets of diversity (adding novel metrics of disparity and variety), diversity dimensions (adding topic diversity and combinations of dimensions) and broader policy context (adding a discussion on existing diversity initiatives and affinity groups). This joint analysis of AAAC, TAFFC and ACII, which covers a decade of community members’ data (from January 2011 to December 2021), firstly allows to reinforce our initial hypothesis on the “leaky pipeline” existing in the field, even though the share of women seems to timidly increase over the years for the three sources. Second, it demonstrates that academic institutions dominate affective computing peer-reviewed publications (both for the journal and the conference), ahead of private industry –which is however increasingly present– and governmental research centres. In terms of countries, USA-based institutions clearly lead the list of publications, while those from Latin America and Africa are vastly under-represented. Lastly, we demonstrate that diversity in the characteristics of led authors (namely their gender and geographic location) has an influence on topic diversity.

To the best of our knowledge, this is the largest and most comprehensive diversity study to date on Affective Computing research. Beyond the Affective Computing field, we provide a novel framework for the quantitative analysis of diversity in a research community, which is grounded on metrics adapted from fields as varied as Economics, Neuroscience or Ecology and that can be used for the analysis of any other community.

Nevertheless, our study has to be considered in light of some limitations. First, note that all connections between different diversity dimensions that we discuss in this work and the intertemporal analysis of topics do not go beyond the analysis of correlations and indicative trends. To explore and confirm drivers, potential confounders and mechanisms of these correlations and trends requires further causal analysis that could exploit an exogenous variation for every hypothesis made throughout this analysis. Nevertheless this work was useful to formulate

these hypotheses in the first place. Moreover, there are other relevant diversity dimensions not covered here. In the context of anti-discrimination and for the purpose of inclusive research and innovation, dimensions such as race, sexual orientation, religion or disability are at least equally important. However, there is no ethical basis for the inference of these dimensions based on name and affiliations. We note that the gender and nationality inference may not be fully adequate either, although we want to highlight the importance of monitoring the diversity of research communities in general. Moreover, gender had to be limited to a binary categorisation and may be misaligned with individual self-identification. Some of these issues could be solved by undertaking diversity initiatives, by explicitly asking community members to self-identify themselves (e.g., in surveys, at registration in ACII), and by fostering collaborations with affinity groups from under-represented minorities.

Although the mentioned limitations reflect the complexity of developing a holistic, robust and universal analysis of diversity in a research community, goal that we will keep on pursuing in our future research, our study identifies specific paths for methodological research and concrete diversity gaps in the affective computing community.

Finally, although we consider diversity an important goal that is valuable in and of its own right, we would also like to provide more support to the claim that “diversity is important” in future work by identifying and highlighting a link between diversity and innovative ideas.

ACKNOWLEDGMENTS

The authors would like to thank the Association for the Advancement of Affective Computing (AAAC) and anonymous reviewers for their support.

REFERENCES

- [1] H. L. E. G. on AI, *Ethics Guidelines for Trustworthy AI*. Brussels, Belgium: European Commission, 2019.
- [2] United Nations, *Universal Declaration of Human Rights*. Paris, France: General Assembly Resolution 217 A, 1948.
- [3] European Union, *Charter of Fundamental Rights of the European Union*. Brussels, Belgium: Official Journal of the European Union, 2012.
- [4] S. M. West, M. Whittaker, and K. Crawford, “Discriminating systems: Gender, race and power in AI,” *AI Now Inst.*, Apr. 2019. [Online]. Available: <https://ainowinstitute.org/discriminatingystems.html>
- [5] R. Picard, *Affective Computing*. Cambridge, MA, USA: The MIT Press, 1997.
- [6] E. Gómez, V. Charisi, S. Tolan, M. Miron, F. Martinez-Plumed, and M. Escobar-Planas, *HUMAIN: Understanding the Impact of Artificial Intelligence on Human Behaviour*. G. Amran, Ed., Luxembourg: Publications Office of the European Union, 2021.
- [7] S. Rukavina, S. Gruss, H. Hoffmann, J.-W. Tan, S. Walter, and H. C. Traue, “Affective computing and the impact of gender and age,” *PLoS One*, vol. 11, no. 3, 2016, Art. no. e0150584.
- [8] Y. Fan, J. C. Lam, and V. O. Li, “Demographic effects on facial emotion expression: An interdisciplinary investigation of the facial action units of happiness,” *Sci. Rep.*, vol. 11, no. 1, pp. 1–11, 2021.
- [9] E. Cambria, Q. Liu, S. Decherchi, F. Xing, and K. Kwok, “SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis,” in *Proc. 13th Conf. Lang. Resour. Eval.*, 2022, pp. 3829–3839.
- [10] A. Khatua, E. Cambria, K. Ghosh, N. Chaki, and A. Khatua, “Tweeting in support of LGBT? A deep learning approach,” in *Proc. ACM India Joint Int. Conf. Data Sci. Manage. Data*, 2019, pp. 342–345.
- [11] A. Stirling, “A general framework for analysing diversity in science, technology and society,” *J. Roy. Soc. Interface*, vol. 4, no. 15, pp. 707–719, 2007.
- [12] R. B. Freeman and W. Huang, “Collaboration: Strength in diversity,” *Nat. News*, vol. 513, no. 7518, 2014, Art. no. 305.
- [13] B. K. AlShebli, T. Rahwan, and W. L. Woon, “The preeminence of ethnic diversity in scientific collaboration,” *Nat. Commun.*, vol. 9, no. 1, pp. 1–10, 2018.
- [14] T. H. Swartz, A.-G. S. Palermo, S. K. Masur, and J. A. Aberg, “The science and value of diversity: Closing the gaps in our understanding of inclusion and diversity,” *J. Infect. Dis.*, vol. 220, pp. S33–S41, 2019.
- [15] J. Adams, “The fourth age of research,” *Nature*, vol. 497, no. 7451, pp. 557–560, 2013.
- [16] I. Hupont, S. Tolan, A. Freire, L. Porcaro, S. Estevez, and E. Gómez, “How diverse is the ACII community? Analysing gender, geographical and business diversity of affective computing research,” in *Proc. 9th Int. Conf. Affect. Comput. Intell. Interaction*, 2021, pp. 1–8.
- [17] I. Hupont, E. Gomez, S. Tolan, L. Porcaro, and A. Freire, “Monitoring diversity of AI conferences: Lessons learnt and future challenges in the DivinAI project,” 2022, *arXiv:2203.01657*.
- [18] A. M. M. Oswald and S. Ostojic, “Curating more diverse scientific conferences,” *Nat. Rev. Neurosci.*, vol. 21, no. 11, pp. 589–590, 2020.
- [19] L. A. Wheaton, “Racial equity and inclusion still lacking in neuroscience meetings,” *Nat. Neurosci.*, vol. 24, no. 12, pp. 1645–1647, 2021.
- [20] A. Vila-Concejo et al., “Steps to improve gender diversity in coastal geoscience and engineering,” *Palgrave Commun.*, vol. 4, no. 1, 2018, Art. no. 103.
- [21] L. King et al., “Diversity in geoscience: Participation, behaviour, and the division of scientific labour at a Canadian geoscience conference,” *Facets*, vol. 3, no. 1, pp. 415–440, 2018.
- [22] K. Dutt, “Race and racism in the geosciences,” *Nat. Geosci.*, vol. 13, no. 1, pp. 2–3, 2020, doi: [10.1038/s41561-019-0519-z](https://doi.org/10.1038/s41561-019-0519-z).
- [23] E. Shishkova, N. W. Kwiecien, A. S. Hebert, M. S. Westphal, J. E. Prenni, and J. J. Coon, “Gender diversity in a STEM subfield – analysis of a large scientific society and its annual conferences,” *J. Amer. Soc. Mass Spectrometry*, vol. 28, no. 12, pp. 2523–2531, 2017.
- [24] N. D. Spector et al., “Women in pediatrics: Progress, barriers, and opportunities for equity, diversity, and inclusion,” *Pediatrics*, vol. 144, no. 5, 2019, Art. no. e20192149.
- [25] B. Partiali, S. Oska, R. B. Touriel, A. Delise, A. Barbat, and A. Folbe, “Gender disparity in speakers at a major academic emergency medicine conference,” *Emerg. Med. J.*, vol. 38, no. 5, pp. 379–380, 2021. [Online]. Available: <https://emj.bmj.com/content/38/5/379>
- [26] W. P. M. Dijksterhuis et al., “From presentation to paper: Gender disparities in oncological research,” *Int. J. Cancer*, vol. 146, no. 11, pp. 3011–3021, jun 2020.
- [27] D. Zhang et al., “2021 AI index report,” Stanford University, 2021. [Online]. Available: <https://aiindex.stanford.edu/report/>
- [28] R. Righi, C. Pineda León, M. Cardona, and J. S. Garrido, “AI watch index 2021,” Publications Office of the European Union, Luxembourg, 2021. [Online]. Available: https://ai-watch.ec.europa.eu/publications/ai-watch-index-2021_en
- [29] A. Freire, L. Porcaro, and E. Gómez, “Measuring diversity of artificial intelligence conferences,” in *Proc. AAAI Workshop Artif. Intell. Diversity Belonging Equity Inclusion*, 2021, pp. 39–50.
- [30] F. Guo, F. Li, W. Lv, L. Liu, and V. G. Duffy, “Bibliometric analysis of affective computing researches during 1999~2018,” *Int. J. Hum.-Comput. Interact.*, vol. 36, no. 9, pp. 801–814, 2020, doi: [10.1080/10447318.2019.1688985](https://doi.org/10.1080/10447318.2019.1688985).
- [31] M.-T. Ho, H.-K. T. Nguyen, M. Peter, and Q.-H. Vuong, “Affective computing at the edge: A bibliometric analysis of the period 1995–2020,” *Emotional AI in Cities: Cross Cultural Lessons from UK and Japan on Designing for an Ethical Life*, 2021. [Online]. Available: <https://osf.io/hu54t/download>
- [32] M. H. Pestana, W.-C. Wang, and L. A. Moutinho, “Global affective computing research in the period 1997–2017: A bibliometric analysis,” *Int. J. Multivariate Data Anal.*, vol. 1, no. 4, pp. 348–370, 2018.
- [33] R. B. Freeman and W. Huang, “Collaborating with people like me: Ethnic coauthorship within the United States,” *J. Labor Econ.*, vol. 33, no. S1, pp. S289–S318, 2015.
- [34] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, “Fast unfolding of communities in large networks,” *J. Statist. Mechanics: Theory Experiment*, vol. 2008, no. 10, 2008, Art. no. P10008.
- [35] J. Daniels, M. Nkonde, and D. Mir, “Advancing racial literacy in tech: Why ethics, diversity in hiring and implicit bias trainings aren’t enough,” *Data Soc. Res. Inst.*, 2019. [Online]. Available: https://img-cache.oppcdn.com/fixed/33022/assets/_YHRDa27Q7vMq4Ku.pdf

- [36] C. Collett and S. Dillon, "AI and gender: Four proposals for future research," *Leverhulme Centre Future Intell.*, Univ. Cambridge, 2019. [Online]. Available: https://www.repository.cam.ac.uk/bitstream/handle/1810/294360/AI_and_Gender_4_Proposals_for_Future_Research_210619_p8qAu8L%20%281%29.pdf
- [37] M. Andrus and S. Villeneuve, "Demographic-reliant algorithmic fairness: Characterizing the risks of demographic data collection in the pursuit of fairness," in *Proc. ACM FAccT Conf.*, 2022. [Online]. Available: <https://arxiv.org/abs/2205.01038>
- [38] The Verge, "Automatic gender recognition tech is dangerous, say campaigners: It's time to ban it," 2021. [Online]. Available: <https://www.theverge.com/2021/4/14/22381370/automatic-gender-recognition-sexual-orientation-facial-ai-analysis-ban-campaign>
- [39] S. Page, "Diversity and complexity," in *Diversity and Complexity*. Princeton, NJ, USA: Princeton Univ. Press, 2010.
- [40] N. Duch-Brown, E. Gomez-Herrera, F. Mueller-Langer, and S. Tolan, "Market power and artificial intelligence work on online labour markets," *Res. Policy*, vol. 51, no. 3, 2022, Art. no. 104446.
- [41] C. E. Shannon, "A mathematical theory of communication," *Bell System Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.
- [42] D. Mouillot and A. Lepretre, "A comparison of species diversity estimators," *Researches Popul. Ecol.*, vol. 41, no. 2, pp. 203–215, 1999.
- [43] J. C. Blickenstaff, "Women and science careers: Leaky pipeline or gender filter?," *Gender Educ.*, vol. 17, no. 4, pp. 369–386, 2005.
- [44] A. Abdelrazek, Y. Eid, E. Gawish, W. Medhat, and A. Hassan, "Topic modeling algorithms and applications: A survey," *Inf. Syst.*, vol. 112, 2022, Art. no. 102131.
- [45] A. Alabdulkareem, M. R. Frank, L. Sun, B. AlShebli, C. Hidalgo, and I. Rahwan, "Unpacking the polarization of workplace skills," *Sci. Adv.*, vol. 4, no. 7, 2018, Art. no. eaao6030.
- [46] O. Sporns and R. F. Betzel, "Modular brain networks," *Annu. Rev. Psychol.*, vol. 67, 2016, Art. no. 613.
- [47] Y. Cohen, D. Hender, and A. Rubin, "Detection of malicious webmail attachments based on propagation patterns," *Knowl.-Based Syst.*, vol. 141, pp. 67–79, 2018.
- [48] T. Cahlik, "Comparison of the maps of science," *Scientometrics*, vol. 49, no. 3, pp. 373–387, 2000.



computing, and face recognition.

Isabelle Hupont received the PhD degree in computer science from the University of Zaragoza, in 2010. She has worked as a research manager with the Aragon Institute of Technology (2006-2015), senior researcher with the Institute of Intelligent Systems and Robotics, Sorbonne University, Paris, 2015-2017, and chief research officer with Herta, Barcelona, 2017-2021. He is now a scientific project officer with the Joint Research Centre of the European Commission, Seville, Spain. Her research interests include focuses on trustworthy artificial intelligence, affective




work is published in several peer-reviewed economics and AI-related journals and conferences.

Songül Tolan received the master's degree in economics from the University of Warwick, and the PhD degree in economics from the Free University of Berlin. Prior to joining the European Commission's Joint Research Centre, she has been a researcher with the Public Policy Department, German Institut for Economic Research (DIW, Berlin). Her research interests include focuses on algorithmic fairness, trustworthy AI, labour market impacts of AI, deploying methods both from applied empirical microeconomics and machine learning. Her inter-disciplinary



a consultant for the European Commission's Joint Research Centre (Seville, Spain) developing tools for AI-related studies.

Pedro Frau received the master's degree in artificial intelligence. He is an audiovisual system engineer. He has been working in several fields related to computer science, mainly in processes automation, aerial robotics and social robotics. He has collaborated with academic groups (e.g., SPECS group with Universitat Pompeu Fabra and CVAR group with the Technical University of Madrid). At the corporate level, he has worked in tech consultancy companies (Everis, NTTDisruption, AKKA technologies) and in defense (Airbus Defense and Space). He is currently



Lorenzo Porcaro received the master's degrees in sound and music computing and intelligent interactive systems from Universitat Pompeu Fabra, Spain, and the PhD degree in information and communication technology. He is currently working as scientific project officer with the European Commission's Joint Research Centre. His research interests include recommender systems, social computing, and human-computer interaction.



company (BMAT). She is currently a member of the Spanish National Council for AI and the OECD One AI expert group.

Emilia Gómez principal investigator on Human and Machine Intelligence (HUMAIN) with the Joint Research Centre, European Commission and guest professor with Universitat Pompeu Fabra, Barcelona. Her research is grounded on the Music Information Retrieval field, where she develops technologies to support music listening experiences. Starting from music, she researches on the ethical and social impact of AI. She has co-authored more than 150 peer-reviewed publications, open datasets and software packages, supervised 11 PhD theses, co-founded a